CLAIMS

1. A method comprising:

transforming frames in a video sequence using a wavelet transform and motion information between frames to produce multiple sub-bands of coefficients; and

coding the coefficients of each sub-band independently.

2. A method as recited in claim 1, wherein the wavelet transform comprises a shape-adaptive discrete wavelet transform.

3. A method as recited in claim 1, wherein the transforming comprises performing a temporal 1-D wavelet transform along motion trajectories in a temporal direction.

4. A method as recited in claim 1, wherein the transforming comprises:

performing a temporal wavelet transform on corresponding pixels in a video object along motion trajectories in a temporal direction to produce frames of temporal wavelet coefficients; and

performing a spatial wavelet transform on the frames of the temporal wavelet coefficients to produce multiple sub-bands of wavelet coefficients.

5. A method as recited in claim 1, wherein the coding produces multiple bitstreams, one for each sub-band, and further comprising forming a bitstream from the multiple bitstreams.

6. A method as recited in claim 1, wherein the coding comprises transposing selected sub-bands.

7. A method as recited in claim 1, wherein the coding comprises coding the coefficients of each sub-band bit-plane by bit-plane using different coding primitives.

8. A method as recited in claim 7, wherein the coding primitives comprise:

zero coding to code new information about a coefficient that is not yet significant in a previous bit-plane; and

sign coding to code a sign of the coefficient once the coefficient is deemed significant.

9. A method as recited in claim 7, wherein the coding primitives comprise:

zero coding to code new information about a coefficient that is not yet significant in a previous bit-plane;

sign coding to code a sign of the coefficient once the coefficient is deemed significant; and

magnitude refinement to code new information of a coefficient that has already become significant in the previous bit-plane.

10. A method as recited in claim 1, wherein the coding comprises assigning contexts to the coefficients of each sub-band based on numbers of significant neighboring samples.

11. A method as recited in claim 10, wherein the sub-bands include an LLL (low-low-low) sub-band and an LLH (low-low-high) sub-band and the contexts are assigned as follows:

LLL and LLH Sub-bands				
h	**v**	**a**	**d**	**Context**
2	x	x	x	0
1	≥1	x	x	0
1	0	≥1	x	1
1	0	0	x	2
0	2	0	x	3
0	1	0	x	4
0	0	≥1	x	5
0	0	0	3	6
0	0	0	2	7
0	0	0	1	8
0	0	0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

12. A method as recited in claim 10, wherein the sub-bands include an LHH (low-high-high) sub-band and the contexts are assigned as follows:

LHH Sub-band			
h	**v+a**	**d**	**Context**
2	x	x	0
1	≥3	x	0
1	≥1	≥4	1
1	≥1	x	2
1	0	≥4	3
1	0	x	4
0	≥3	x	5
0	≥1	≥4	6
0	≥1	x	7
0	0	≥4	8
0	0	x	9

where “h” represents a number of immediate horizontal neighbors that are significant and $0<h<2$, “v” represents a number of immediate vertical neighbors that are significant and $0<v<2$, “a” represents a number of immediate temporal neighbors that are significant and $0<a<2$, and “d” represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

13. A method as recited in claim 10, wherein the sub-bands include an HHH (high-high-high) sub-band and the contexts are assigned as follows:

d	h+v+a	Context
≥6	x	0
≥4	≥3	1
≥4	x	2
≥2	≥4	3
≥2	≥2	4
≥2	x	5
≥0	≥4	6
≥0	≥2	7
≥0	1	8
≥0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

14. A method as recited in claim 1, further comprising truncating a number of bits in each bit-plane according to rate-distortion curves.

15. A method as recited in claim 1, further comprising estimating motion trajectories of pixels in a video object from frame to frame in the video sequence and said transforming is performed on corresponding pixels along the motion trajectories.

16. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 1.

17. A method comprising:

estimating motion trajectories of pixels in a video object from frame to frame in a video sequence;

performing a temporal wavelet transform on the corresponding pixels along the motion trajectories in a temporal direction to produce frames of temporal wavelet coefficients;

performing a spatial wavelet transform on the frames of the temporal wavelet coefficients to produce multiple sub-bands of wavelet coefficients; and

coding each sub-band of wavelet coefficients independently.

18. A method as recited in claim 17, wherein the estimating comprises matching corresponding pixels in the video object from frame to frame in the video sequence.

19. A method as recited in claim 17, wherein the temporal and spatial wavelet transforms comprise a shape-adaptive discrete wavelet transform.

20. A method as recited in claim 17, wherein the performing a temporal wavelet transform comprises:

forming a pixel array containing pixels that continue from frame to frame in the video sequence;

examining a pixel in a frame to determine whether the pixel is a terminating pixel that does not continue to a next frame;

if the pixel is a terminating pixel, terminating the pixel array; and

if the pixel is not a terminating pixel, adding the pixel to the pixel array.

21. A method as recited in claim 20, further comprising transforming the pixels arrays to produce the frames of temporal wavelet coefficients.

22. A method as recited in claim 17, wherein the coding comprises transposing selected sub-bands to reduce a number of sub-bands to be coded.

23. A method as recited in claim 17, wherein the coding comprises:

coding the wavelet coefficients in bit-planes; and

allocating bits for the bit-planes according to a rate-distortion optimization.

24. A method as recited in claim 17, further comprising truncating bits allocated to a bit-plane at a point on a rate-distortion curve that approximates a convex hull.

25. A method as recited in claim 17, wherein the coding comprises coding the wavelet coefficients of each sub-band bit-plane by bit-plane using different coding primitives.

26. A method as recited in claim 25, wherein the coding primitives comprise:

zero coding to code new information about a wavelet coefficient that is not yet significant in a previous bit-plane; and

sign coding to code a sign of the wavelet coefficient once the wavelet coefficient is deemed significant.

27. A method as recited in claim 25, wherein the coding primitives comprise:

zero coding to code new information about a wavelet coefficient that is not yet significant in a previous bit-plane;

sign coding to code a sign of the wavelet coefficient once the wavelet coefficient is deemed significant; and

magnitude refinement to code new information of a wavelet coefficient that has already become significant in the previous bit-plane.

28. A method as recited in claim 17, wherein the coding produces multiple bitstreams for corresponding sub-bands of wavelet coefficients and further comprising constructing a multi-layer bitstream from the multiple bitstreams.

29. A method as recited in claim 17, wherein the coding comprises assigning contexts to the wavelet coefficients of each sub-band based on numbers of significant neighboring samples.

30. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 17.

31. A method comprising:

forming an array containing a current pixel in a current frame;

determining whether a pixel corresponding to the current pixel resides in a next frame subsequent to the current frame;

if the corresponding pixel exists, add the current pixel to the array; and

if the corresponding pixel does not exist, terminate the pixel array with the current pixel.

32. A method as recited in claim 31, further comprising initially indicating all pixels as one state and changing the state of each pixel after said examining of the pixel.

33. A method as recited in claim 31, further comprising applying a wavelet transform to the pixel array to produce wavelet coefficients.

34. A method as recited in claim 31, further comprising:

applying a wavelet transform to the pixel array to produce wavelet coefficients; and

applying a spatial 2-D wavelet transform on the wavelet coefficients.

35. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 31.

36. A method comprising:

coding sub-bands of coefficients produced from transforming video frames in an independent manner such that one sub-band of coefficients is coded independently of another sub-band of coefficients; and

constructing a bitstream from the independently coded sub-bands.

37. A method as recited in claim 36, wherein the coding comprises transposing selected sub-bands prior to said coding.

38. A method as recited in claim 36, wherein the coding comprises coding the coefficients of each sub-band bit-plane by bit-plane using different coding primitives.

39. A method as recited in claim 38, wherein the coding primitives comprise:

zero coding to code new information about a coefficient that is not yet significant in a previous bit-plane; and

sign coding to code a sign of the coefficient once the coefficient is deemed significant.

40. A method as recited in claim 38, wherein the coding primitives comprise:

zero coding to code new information about a coefficient that is not yet significant in a previous bit-plane;

sign coding to code a sign of the coefficient once the coefficient is deemed significant; and

magnitude refinement to code new information of a coefficient that has already become significant in the previous bit-plane.

41. A method as recited in claim 36, wherein the coding comprises assigning contexts to the coefficients of each sub-band based on numbers of significant neighboring samples.

42. A method as recited in claim 41, wherein the sub-bands include an LLL (low-low-low) sub-band and an LLH (low-low-high) sub-band and the contexts are assigned as follows:

LLL and LLH Sub-bands				
h	**v**	**a**	**d**	**Context**
2	x	x	x	0
1	≥1	x	x	0
1	0	≥1	x	1
1	0	0	x	2
0	2	0	x	3
0	1	0	x	4
0	0	≥1	x	5
0	0	0	3	6
0	0	0	2	7
0	0	0	1	8
0	0	0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

43. A method as recited in claim 41, wherein the sub-bands include an LHH (low-high-high) sub-band and the contexts are assigned as follows:

LHH Sub-band			
h	**v+a**	**d**	**Context**
2	x	x	0
1	≥3	x	0
1	≥1	≥4	1
1	≥1	x	2
1	0	≥4	3
1	0	x	4
0	≥3	x	5
0	≥1	≥4	6
0	≥1	x	7
0	0	≥4	8
0	0	x	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

44. A method as recited in claim 41, wherein the sub-bands include an HHH (high-high-high) sub-band and the contexts are assigned as follows:

d	h+v+a	Context
≥6	x	0
≥4	≥3	1
≥4	x	2
≥2	≥4	3
≥2	≥2	4
≥2	x	5
≥0	≥4	6
≥0	≥2	7
≥0	1	8
≥0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

45. A method as recited in claim 36, wherein the constructing comprises forming multiple bit-planes and truncating a number of bits in each bit-plane according to a rate-distortion curve.

46. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 36.

47. A method comprising:

transforming a video sequence to produce a set of sub-bands;

transposing selected sub-bands to produce a reduced set of sub-bands that are fewer than the set of sub-bands; and

coding the reduced set of sub-bands.

48. A method as recited in claim 47, wherein the coding comprises coding each sub-band independently.

49. A method as recited in claim 47, wherein the coding comprises assigning contexts to the coefficients of each sub-band based on numbers of significant neighboring samples.

50. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 47.

51. A method for coding coefficients indicative of transformed video in multiple bit-planes, comprising:

conducting coding passes through a bit-plane to code significant samples separately from insignificant samples; and

repeating said conducting for each bit-plane.

52. A method as recited in claim 51, wherein the three coding passes comprise:

a significant propagation pass to code samples that are not yet significant;

a magnitude refinement pass to code samples that are already deemed significant; and

a normalization pass to code any samples not coded in the significant propagation pass and the magnitude refinement passes.

53. A computer-readable medium comprising computer-executable instructions that, when executed by one or more processors, perform the method as recited in claim 51.

54. A video encoder comprising:

a wavelet transformer to transform frames in a video sequence into multiple sub-bands of coefficients, the wavelet transform using motion information of video objects in the frames; and

a coder to code the coefficients of each sub-band independently.

55. A video encoder as recited in claim 54, wherein the wavelet transformer applies a spatial-adaptive discrete wavelet transform to the frames.

56. A video encoder as recited in claim 54, wherein the wavelet transformer applies a temporal wavelet transform on pixels in the frames taken along motion trajectories in a temporal direction from frame to frame.

57. A video encoder as recited in claim 54, wherein the wavelet transformer comprises a 3-D wavelet transformer that applies:

(1) a temporal 1-D wavelet transform on corresponding pixels in consecutive frames along motion trajectories in a temporal direction to produce temporal wavelet coefficients; and

(2) a spatial 2-D wavelet transform on the temporal wavelet coefficients.

58. A video encoder as recited in claim 54, wherein the wavelet transformer estimates motion trajectories of pixels in a video object from frame to frame in the video sequence and initially transforms corresponding pixels along the motion trajectories in the temporal direction.

59. A video encoder as recited in claim 54, wherein the coder codes transposes selected sub-bands to produced reduced set of sub-bands.

60. A video encoder as recited in claim 54, wherein the coder codes the coefficients of each sub-band into bit-planes using different coding primitives.

61. A video encoder as recited in claim 54, wherein the coder comprises a context-based arithmetic coder to assign contexts to the coefficients of each sub-band based on different coding primitives.

62. A video encoder as recited in claim 61, wherein the sub-bands include an LLL (low-low-low) sub-band and an LLH (low-low-high) sub-band and the coder employs a zero coding primitive to code new information about a coefficient that is not yet significant in a previous bit-plane by assigning the contexts as follows:

LLL and LLH Sub-bands				
h	**v**	**a**	**d**	**Context**
2	x	x	x	0
1	≥1	x	x	0
1	0	≥1	x	1
1	0	0	x	2
0	2	0	x	3
0	1	0	x	4
0	0	≥1	x	5
0	0	0	3	6
0	0	0	2	7
0	0	0	1	8
0	0	0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

63. A video encoder as recited in claim 61, wherein the sub-bands include an LHH (low-high-high) sub-band and the coder employs a zero coding primitive to code new information about a coefficient that is not yet significant in a previous bit-plane by assigning the contexts as follows:

LHH Sub-band			
h	**v+a**	**d**	**Context**
2	x	x	0
1	≥3	x	0
1	≥1	≥4	1
1	≥1	x	2
1	0	≥4	3
1	0	x	4
0	≥3	x	5
0	≥1	≥4	6
0	≥1	x	7
0	0	≥4	8
0	0	x	9

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and "d" represents a number of immediate diagonal neighbors that are significant and $0<d<12$.

64. A video encoder as recited in claim 61, wherein the sub-bands include an HHH (high-high-high) sub-band and the coder employs a zero coding primitive to code new information about a coefficient that is not yet significant in a previous bit-plane by assigning the contexts as follows:

d	h+v+a	Context
≥6	x	0
≥4	≥3	1
≥4	x	2
≥2	≥4	3
≥2	≥2	4
≥2	x	5
≥0	≥4	6
≥0	≥2	7
≥0	1	8
≥0	0	9

where "h" represents a number of immediate horizontal neighbors that are significant and 0<h<2, "v" represents a number of immediate vertical neighbors that are significant and 0<v<2, "a" represents a number of immediate temporal neighbors that are significant and 0<a<2, and "d" represents a number of immediate diagonal neighbors that are significant and 0<d<12.

65. A video encoder as recited in claim 61, wherein the coder employs a sign coding primitive to code a sign of the coefficient once the coefficient is deemed significant by assigning the contexts as follows:

h=-1			
v	a	$\hat{\chi}$	Context
-1	-1	0	0
-1	0	0	1
-1	1	0	2
0	-1	0	3
0	0	0	4
0	1	0	5
1	-1	0	6
1	0	0	7
1	1	0	8

H=0			
v	a	$\hat{\chi}$	Context
-1	-1	0	9
-1	0	0	10
-1	1	0	11
0	-1	0	12
0	0	0	13
0	1	1	12
1	-1	1	11
1	0	1	10
1	1	1	9

h=1			
v	a	$\hat{\chi}$	Context
-1	-1	1	8
-1	0	1	7
-1	1	1	6
0	-1	1	5
0	0	1	4
0	1	1	3
1	-1	1	2
1	0	1	1
1	1	1	0

where "h" represents a number of immediate horizontal neighbors that are significant and $0<h<2$, "v" represents a number of immediate vertical neighbors that are significant and $0<v<2$, "a" represents a number of immediate temporal neighbors that are significant and $0<a<2$, and $\hat{\chi}$ is a sign symbol prediction in a given context.

66. A video encoder as recited in claim 54, wherein the coder truncates a number of coding bits according to rate-distortion curves.

67. An operating system embodied on a computer-readable medium comprising a video encoder as recited in claim 54.

68. A video encoder comprising:

means for estimating motion trajectories of pixels in a video object from frame to frame in a video sequence;

means for performing a temporal wavelet transform on the corresponding pixels along the motion trajectories in a temporal direction to produce frames of temporal wavelet coefficients;

means for performing a spatial wavelet transform on the frames of the temporal wavelet coefficients to produce multiple sub-bands of wavelet coefficients; and

means for coding each sub-band of wavelet coefficients independently.

69. A video encoder as recited in claim 68, wherein the estimating means comprises means for matching corresponding pixels in the video object from frame to frame in the video sequence.

70. A video encoder as recited in claim 68, wherein the temporal and spatial wavelet transforms comprise a shape-adaptive discrete wavelet transform.

71. A video encoder as recited in claim 68, wherein the means for performing a temporal wavelet transform comprises:

means for forming a pixel array containing pixels that continue from frame to frame in the video sequence;

means for examining a pixel in a frame to determine whether the pixel is a terminating pixel that does not continue to a next frame;

if the pixel is a terminating pixel, means for terminating the pixel array; and

if the pixel is not a terminating pixel, means for adding the pixel to the pixel array.

72. A video encoder as recited in claim 68, wherein the coding means comprises means for transposing selected sub-bands to reduce a number of sub-bands to be coded.

73. A video encoder as recited in claim 68, wherein the coding means comprises:

means for coding the wavelet coefficients in bit-planes; and

means for allocating bits for the bit-planes according to a rate-distortion optimization.

74. A video encoder as recited in claim 68, further comprising means for truncating bits allocated to a bit-plane at a point on a rate-distortion curve that approximates a convex hull.

75. A video encoder as recited in claim 68, wherein the coding means comprises means for coding the wavelet coefficients of each sub-band bit-plane by bit-plane using different coding primitives.

76. A video encoder as recited in claim 68, wherein the coding means produces multiple bitstreams for corresponding sub-bands of wavelet coefficients and further comprising means for constructing a multi-layer bitstream from the multiple bitstreams.

77. A video encoder as recited in claim 68, wherein the coding means comprises means for assigning contexts to the wavelet coefficients of each sub-band based on numbers of significant neighboring samples.

78. A computer-readable medium comprising computer-executable instructions that, when executed on a processor, direct a device to:

code sub-bands of coefficients produced from transforming video frames in an independent manner such that one sub-band of coefficients is coded independently of another sub-band of coefficients; and

construct a bitstream from the independently coded sub-bands.

79. A computer-readable medium as recited in claim 78, further comprising computer-executable instructions that, when executed on a processor, direct a device to transpose selected sub-bands prior to said coding.

80. A computer-readable medium as recited in claim 78, further comprising computer-executable instructions that, when executed on a processor, direct a device to code the coefficients of each sub-band bit-plane by bit-plane using different coding primitives.

81. A computer-readable medium as recited in claim 78, further comprising computer-executable instructions that, when executed on a processor, direct a device to assign contexts to the coefficients of each sub-band based on numbers of significant neighboring samples.

82. A computer-readable medium embodying an encoded video signal constructed as a result of a process comprising:

transforming frames in a video sequence using a wavelet transform to produce multiple sub-bands of coefficients;

coding the coefficients of each sub-band independently to produce multiple bitstreams; and

forming a bitstream from the multiple bitstreams.